The Cloudlet Bazaar
Dynamic Markets for the Small Cloud

Ranjan Pal, *Member, IEEE,* Sung-Han Lin, *Student Member, IEEE,* Leana Golubchik, *Member, IEEE*

Abstract—Cloud computing is becoming increasingly pervasive due to attractive features such as on-demand resource provisioning and elasticity that it provides. Most commercial cloud providers (e.g., Amazon) are generally centralized entities employing massive data centers that store and process data from various parts of the globe. However, due to increasing concerns about privacy and data control as one of the factors, many small clouds (SCs) established by different providers are emerging in an attempt to meet demand locally [1]. The flip-side of this approach is the problem of resource inelasticity faced by the SCs due to their relatively scarce resources (e.g., virtual machines), thereby leading to potential degradation of customer QoS and loss of revenue. A proposed solution to this problem recommends the sharing of resources between SCs to alleviate the resource inelasticity issues that might arise [1]. However, effective borrowing of resources by an SC from its peers involves mutually satisfying the interests of the stakeholders in question, i.e., the SC customers, the SCs, and a regulatory agency (e.g., the federal government) overseeing the functioning of the SCs. In this paper, we model the 'stakeholder satisfaction problem' in SCs as a *socially efficient* dynamic market/ecosystem design task, where the market elements comprise the stakeholders, and the term 'social efficiency' implying the market reaching a *utilitarian maximum social welfare* state when in equilibrium. The dynamicity of SC markets arises due to the non-static nature of the supply of SC resources, as well as due to the variations in customer demand over time, potentially leading to market disequilibrium. Our market design approach is based on Arrow and Hurwicz's *disequilibrium process* [2, 3], and uses the *gradient play* technique in game theory to iteratively converge upon efficient and stable market equilibria. We illustrate the stability and sustainability of SC markets via both theory and experiments.

***Keywords*—** small cloud, resource sharing, dynamic market, market equilibrium, gradient play, stability

arXiv:1704.00845v1 [cs.DC] 4 Apr 2017

1 Introduction

Cloud computing is becoming increasingly popular and pervasive in the information technology (IT) marketplace due to its on-demand resource provisioning, high availability, and elasticity. These features allow cloud end users (e.g., individuals, small-scale companies, world-wide enterprises) to access resources in a pay-as-you-go manner and to meet varying demands sans upfront resource commitments [4].

1.1 The Rise of Small Clouds

Despite the presence of a number of big cloud providers (e.g., Amazon, Google, Microsoft), many *cloudlets* or small clouds (SCs) [1] are currently springing up in order to avoid severe issues (e.g., legal, and privacy [5] as explained below) that can arise with the adoption of massive centralized data centers capable of storing and processing data from various parts of the globe. Here, the term 'SC' denotes any cloud that is not massive. In addition to legal and privacy issues, massive clouds (MCs) sometimes prove costly and speed-inefficient for certain businesses [6]. As an example, in multi-tenant cloud environments provided by the MCs, it is likely that the latter will find it costly to customize their VMs, and provide subsequent IT support for a large group of their customers, and so will tend to provide one instance satisfying all customers. In comparison, SCs having fewer customers can relatively more cost effectively customize their customer VMs and provide better IT support, in turn attracting more clients. Below, we briefly describe legal and privacy issues that motivate the rise of small clouds. *We emphasize here that we do not mathematically model privacy and/or legal aspects in our work. They just serve as important motivating factors to the rise of small clouds.*

R.Pal is with the jointly with the departments of Computer Science and Electrical Engineering, University of Southern California, CA, USA. E-mail: rpal@usc.edu.
S-H. Lin is with the Department of Computer Science, University of Southern California. E-mail: sunghan@usc.edu.
L. Golubchik is jointly with the departments of Computer Science and Electrical Engineering, University of Southern California. E-mail: leana@usc.edu.

Privacy Issues. By offering a wide variety of services, a big public cloud hosted by a large corporate organization such as Google is able to obtain different pieces of user information which pertain to different fields of endeavor. When users search for something on the web, Google can learn about their interests; when users read their emails on Gmail or check location on Google Maps, Google can learn more about their personal or professional life; when users use the Google Compute Engine, Google can learn about their input data. The greater the scope of the cloud, the greater is the amount of data that can be gathered together and the more valuable is the information that can be obtained with the processing and correlation of such data. While this is likely to help Google increase its profit, the collection and processing of user data into a common integrated framework can also benefit the users when it comes to increasing the quality of the service. Many users are therefore not merely agreeing, but even eager to share their personal data and information with Google in order to obtain a more customized and integrated service. The problem arises when the information given to separate (and apparently independent) services is actually aggregated together by one single entity (either because it is the common provider of said services, or because it has acquired the data from third parties). Even though the information had been voluntarily provided by users, aggregated data might provide further information about users, which they did not necessarily want to disclose.

Legal and Data Control Issues. Exporting data to the cloud means that users can no longer exercise any kind of control over the use and the exploitation of data. Data stored in various small cloud owned data centers can be processed without the knowledge of users, to be further redistributed to third parties without their consent. In addition, the international character of the cloud adds an additional layer of complexity. Information stored in the cloud can be subject to a variety of different laws according to the location where it is being stored or transmitted. A cloud provider might avail itself of the services of other cloud providers located in different jurisdictions, or, if the cloud provider has data centers in multiple countries, it can transfer data between centers depending on economic factors such as

the price of electricity. This means that a file being served from Luxembourg one moment could be served from the Philippines the next. The difficulty for the user to know with certainty which law applies to the information stored into the cloud raises a number of data protection questions.

Emerging SCs. According to recent news reports [7], many European countries are now concerned about their data leaving European borders and are also stressing on the need for local small clouds to serve local needs. For instance, in a country like Switzerland with various autonomous cantons, one could imagine the emergence of multiple small clouds in various government institutions to obviate the reliance on big enterprise data centers and to maintain control over local data [1]. Leaving aside privacy and data control issues, when it comes to servicing future IT demands, by providing space from 10 to hundreds of racks, a small/medium sized data center can (i) serve the requirements of many commercial organizations, government departments, hospitals, universities, etc., and generate revenues, (ii) save operational and energy costs of certain enterprises who instead of owning their own data centers might offload their tasks to SCs, and (iii) generate revenue via outsourcing spare resources to other SCs when the latter run out of resources during periods of workload spikes. To service IT demands, one might even think of having a large data center achieving (i), (ii), and (iii). However, big/massive clouds are often more expensive than SCs, and do not generally provide customized services to their clients.

Despite the potential emergence of small clouds, the latter due to their moderate sizes, are likely to suffer from resource under-provisioning, thus failing to meet peak demand at times. This leads to a resource provisioning dilemma where the SCs have to make the tradeoff between request loss and the cost of over-provisioning. One way out of this dilemma is for such small data centers to cooperate with each other to help meet each others' user demand via resource sharing at low costs, thereby increasing their individual resources when in need without having to significantly invest in more. Such cooperation is analogous to *Business Clusters* described in mainstream economics which emerge due to, among other factors, shared interests and geographical proximity [8]. We note here that since one of the key benefits of SCs is the avoidance of legal and privacy issues, it is apparent that SCs much like Business Clusters would prefer to cooperate with those SCs which lie under one legislative control (modeled in Section 2.3 as a regulatory agency), thereby alleviating any of their own or their customers' privacy concerns.

1.2 Research Motivation

In this section, we briefly describe the problem setting followed by the challenges that motivate us to alleviate them.

Problem Setting. The effective sharing or borrowing resources by an SC from its peers involves mutually satisfying the interests of the stakeholders in context. In this paper, we consider three different stakeholders: (i) the SC customers, (ii) profit maximizing autonomous SCs, and (iii) a regulatory agency overseeing certain functioning aspects of the autonomous SCs (e.g., ensuring customer data privacy). The SC customers are interested in achieving certain performance measures for their jobs (e.g., low job response time, cheap storage); the SCs are interested in maximizing revenues obtained from serving customers; and a regulatory agency (e.g., the local government, a federated agency [9][10]) is interested in (i) ensuring a proper manner by which the autonomous SCs conduct their business of lending resources to peer SCs (e.g., preserving data privacy, designing policies for (a) customer cost effectiveness by disallowing SCs to charge high customer prices, and (b) maintaining a certain level SC and customer welfare), and (ii) recommending resource exchanges between autonomous SCs in a way (without interfering in the important resource (allocation, scheduling) decisions of the SCs) that encompasses the necessary means for the proper configuration of the resources (e.g., using the *OpenNebula* manager and its external resource lease manager Haizea [11]). *We term the above setting as an SC market.*

The Challenges. Ideally, an SC would want to service all its customers solely using its own resources. However, the primary barrier to this goal is its individual resource capacity which might not be enough to service peak customer demand. In such a case, the SC can either resort to peer SCs for additional resources, thereby incurring borrowing costs, and/or buy the services of a big data center. The latter option is generally more expensive than the former and also likely to be more privacy threatening. Thus, from an SC's viewpoint, it's challenge is to satisfy two conflicting objectives: (i) to generate as much revenue by serving its customer demands, and (ii) to incur as low as possible, borrowing and/or buying costs from other data centers. For simplicity purposes, we assume that buying resources from big data centers is the last resort for an SC in events of low resource availability, and in such events it would try its best to get resources from peer SCs. Another challenge is to ensure that at market equilibrium (see below), the SCs and their customers ideally operate on parameters (see Section 2) that allow the market to be *efficient*, a condition commonly characterized in microeconomics by certain popular functions (see Section 2.3) of market stakeholder utilities, and one that entails optimal social welfare allocation amongst the SCs and their customers. This is a non-trivial and challenging task as the existence of a market equilibrium does not necessarily imply market efficiency [12]. In addition to the above mentioned challenges, the SC market is dynamic in nature due to the non-static nature of the supply of SC resources, as well as due to the variations in customer demand over time. This dynamic nature of the SC market is likely to lead to frequent market equilibrium perturbations and potentially a state of market disequilibrium. Here, the term 'equilibrium' refers to a situation in which all market stakeholders mutually satisfy their interests, in which case an important challenge is to design a stable market that is robust to perturbations and always returns to its equilibrium point(s) when perturbations do occur.

Our Goal. *In this paper, our goal is to formulate the joint 'stakeholder satisfaction problem' in SC environments as an efficient, stable, and sustainable dynamic market/ecosystem design task, and propose an effective solution for it.*

1.3 Research Contributions

We make the following research contributions in this paper.

- We propose a utility theory based small cloud competitive market model comprising of SC customers, profit maximizing autonomous SCs, and a regulatory agency overseeing some functionality aspects of the SCs, as the market stakeholders. The model mathematically expresses the stakeholder interests in terms of utility functions and paves the path for analyzing SC markets for market equilibrium properties (see Section 2).
- We analyze our proposed market model via a convex optimization framework for the existence and uniqueness of a static market equilibrium at which (i) the *utilitarian social welfare function* (see Section 3 for a definition) is maximized, i.e., the market equilibrium is socially efficient, (ii) the market equilibrium is *Pareto efficient* (see Section 3 for a definition), (iii) the market is cleared, i.e., the SC supply balances customer demand, and (iv) no stakeholder has any incentive to deviate from the

equilibrium. We show that there exists a unique static competitive market equilibrium jointly satisfying (i), (ii), (iii), and (iv), however there are several static market equilibria jointly satisfying (ii), (iii), and (iv). (see Section 3).

- Using notion of a *disequilibrium process* proposed by Arrow and Hurwicz [2, 3], we then apply the *gradient play* technique in game theory [13] that is based on the theory of differential equations, to investigate the dynamic market setting where a static equilibrium is potentially subject to perturbations that might lead to market disequilibrium. In this regard, we show (in theory) that static market equilibrium achieved in small data center markets is *asymptotically stable* in dynamic market settings. Our use of the gradient play technique is motivated by the fact that in many practical market environments stakeholders (i) find it behaviorally difficult or computationally expensive to play their *best responses* [14], (ii) have zero or incomplete knowledge of the utilities of other stakeholders in the market, and (iii) cannot even observe the actions of other stakeholders in the worst case. In such environments, gradient play is a suitable technique to achieve static market equilibrium stability iteratively [15], from a state of disequilibrium. More specifically, for our market setting the occurrence of (i)-(iii) is quite likely. Gradient play also works when issues (i)-(iii) do not arise (see Section 4).
- We validate our proposed theory through extensive experiments conducted on a synthetic dataset to illustrate the stability of SC markets and the high speed with which such markets converge to stable equilibria despite variations in market supply-demand. Through numerical experiments, we also investigate market equilibria performance of SC markets with respect to three *Bergson-Samuelson* social welfare functions, viz., the *utilitarian function*, the *egalitarian function*, and the *Rawl's function*, standard in the economics literature [12] (see Section 2.3). Here, we study and compare (i) the fairness amongst market players of a welfare allocation achieved through a given social welfare function and (ii) the amount of social welfare of SC markets under the different social welfare functions, when social welfare optimality is achieved. Via our numerical evaluation, we infer that SC markets are sustainable. (see Section 5).

2 Competitive Market Model

In this section, we propose a utility theory based small data center *Walrasian* competitive market model comprising of profit maximizing autonomous SCs, their customers, and a regulatory agency overseeing some functionality aspects of the SCs. A Walrasian competitive market [12] represents a *pure exchange economy* without production, where there are a finite number of agents, i.e., SCs in our work, endowed with a finite number of commodities, i.e., computing resources in our work, that gets traded with SC customers and peer SCs. The aim behind proposing the model is to pave the path for mathematically analyzing SC markets for market equilibrium properties, and derive their practical implications.

In this paper, we consider each SC customer to deal with three job types, where each job comprises multiple tasks: (i) Type I jobs that need to be serviced *wholly/entirely* when they arrive (e.g., a user could invoke a regular MapReduce batch job that defines a set of Mappers and Reducers to be executed for the job to complete in its entirety.), (ii) Type II jobs that can be *curtailed* to fewer tasks (e.g., an approximation job like the ones cited in [16]), the curtailment decision primarily arising from (a) the nature of VM instance prices, (b) the unnecessity of the customer to keep executing a job beyond a certain accuracy already achieved, and (c) the unnecessity of the customer to keep executing the job beyond a certain deadline, and (iii) Type III jobs where certain tasks can be *shifted* over time for future processing, the remaining job tasks requiring service as they arrive (e.g., analyzing a DNA sequence, re-running partially/entirely a current job later when it gets killed in a spot cloud environment due to momentary unavailability of resources.). Next, we model the stakeholders in the SC market.

2.1 Modeling the SCs

Let there be n autonomous profit maximizing SCs. Each SC can spread over multiple locations. Customer demand for SC i is a set of processing tasks from its customers (both end-users and peer SCs) that require the use of virtual machines as the primary computing resources. In this regard, we assume that each SC i reserves (allocates) a total of vm_i^r virtual machines (VMs) in its data centers to service demands from its customers. We term such VMs as *reserved* VMs. The value of vm_i^r is predecided by SC i based on the statistics of customer demand patterns observed over a period of time, as one of the factors. For simplicity, we will focus on VMs representing a single resource type in this paper. This assumption can be easily relaxed by adopting a more elementary resource measure such as Amazon EC2 compute units, which are used as building blocks for different types of VMs. In the event that vm_i^r machines are insufficient to satisfy consumer demands, SC i buys/borrows vm_i^b VMs from peer SCs. Here, vm_i^b is the number of *borrowed* VMs available to SC i from its peers. In the event that both reserved and borrowed VMs are insufficient to meet customer demand, SC i resorts to a *public cloud* for vm_i^{pc} VM instances. We assume here that a public cloud is large enough to provide any required number of VM instances to SCs. We do not consider communication network bandwidth issues to be a bottleneck to customer service satisfaction in our work.

Let $c(vm_i^r)$ be the associated operating cost to SC i for reserving vm_i^r of its own virtual machines to serve its customers (includes peer SCs who borrow resources from i). We define $c(vm_i^r)$ via a separable equation of the following form.

$$c(vm_i^r) = f_1(vm_i^r) + f_2(vm_i^r), \tag{1}$$

where $f_1(\cdot)$ (a linear function) and $f_1(\cdot)$ (a non-linear function) are functions such that the *marginal operating cost* for SC i is a *general decreasing linear function* of the number of VM instances, i.e., the additional operating cost, $\frac{dc}{dvm_i^r}$, due to a unit increase in the number of VMs required to service customer demand varies in a negative linear fashion with the number of VMs. We use this type of marginal cost functions in our work due to their popularity in economics to model diminishing costs/returns [12]. We approximate the number of VMs to be a non-discrete quantity. Specifically, for the purpose of analysis, we assume the cost function $c(\cdot)$ to be *concave*, *quadratic*, and *twice continuously differentiable*, i.e., the marginal costs become decreasing linear functions of the number of VM instances. We can define one such $c(vm_i^r)$ function as follows.

$$c(vm_i^r) = \alpha_r^i vm_i^r + \frac{\beta_r^i}{2}(vm_i^r)^2, \tag{2}$$

where α_r^i (a positive value) and β_r^i (a negative value) are SC i's cost coefficients for its reserved resources, i.e., virtual machines, such that the marginal operating cost for SC i is a negative linear function. *The above quadratic form of the cost function, apart from satisfying the property of negative linear marginals, not only allows for tractable analysis, but also serves as a good second-order approximation for the broader class of concave payoffs [17].* We denote π_i^r to be the profit that SC i makes through its reserved

VMs for servicing customers, and define the maximum profit that SC i can make, via the following optimization problem.

$$\max_{vm_i^r} \pi_i^r = \max_{vm_i^r}[\rho_i vm_i^r - c(vm_i^r)]$$

subject to

$$vm_{\min_i}^r \leq vm_i^r \leq vm_{\max_i}^r,$$

where ρ_i is the per-unit VM instance price charged by SC i to its customers, and $vm_{\min_i}^r$ and $vm_{\max_i}^r$ are the lower and upper bounds for the number of VM instances reserved by SC i for its customers. We assume that each SC i is small enough not to be able to exert market power over its peer SCs and strategically influence the prices they charge their customers. i.e., each SC is a *price taker* [12]. The prices that individual SCs charge their customers are determined by either (i) a regulatory agency aiming to achieve a certain optimal social welfare state (see Section 2.3) or (ii) individual SCs in price competition with one another in the process of maximizing their own utilities and selling off their endowment. In an autonomous setting, (ii) is more likely.

Let $c(vm_i^b)$ be the associated operating cost to SC i for borrowing vm_i^b virtual machines from peer SCs to serve customers, when the reserved VMs are not enough to satisfy customer service demands. The operating cost includes the cost (including transferring data) to push work to the peer SCs from which i borrows resources. Like in the case of formulating $c(vm_i^r)$, we formulate $c(vm_i^b)$ in a manner such that the associated marginal operating costs for borrowing an additional VM instance decreases in a negative linear fashion with the number of VMs. Mathematically, we represent $c(vm_i^b)$ by the following equation.

$$c(vm_i^b) = \alpha_b^i vm_i^b + \frac{\beta_b^i}{2}(vm_i^b)^2, \tag{3}$$

where α_b^i (a positive quantity) and β_b^i (a negative quantity) are SC i's coefficients for its borrowed virtual machines. We denote by π_i^b the profit that SC i makes when borrowing VMs from peer SCs for servicing customers, and define the maximum profit that SC i can make, via the following optimization problem.

$$\max_{vm_i^b} \pi_i^b = \max_{vm_i^b}[\rho_i vm_i^b - c(vm_i^b) - c(vm_i^{pc})]$$

subject to

$$vm_{\min_i}^b \leq vm_i^b \leq vm_{\max_i}^b.$$

Here, (i) $vm_{\min_i}^b$ and $vm_{\max_i}^b$ are the lower and upper bounds for the number of VM instances borrowed by SC i for its customers, from peer SCs, (ii) $c(vm_i^{pc})$ is the cost to SC i to offload vm_i^{pc} VM instances worth of customer demand to a public cloud in the event that vm_i^r and vm_i^b VM instances together are not enough to service i's total customer demand. We mathematically represent $c(vm_i^{pc})$ in the same manner as $c(vm_i^r)$ and $c(vm_i^b)$, and express it via the following equation.

$$c(vm_i^{pc}) = \alpha_{pc}^i vm_i^{pc} + \frac{\beta_{pc}^i}{2}(vm_i^{pc})^2, \tag{4}$$

where α_{pc}^i (a positive quantity) and β_{pc}^i (a negative quantity) are SC i's coefficients for the resources the public cloud uses to service i's offloaded customer demand portions. We do not assume any constraints on the resources available to the public cloud for servicing offloading requests by SCs. Next, we model SC customers.

2.2 Modeling SC Customers

For a customer j who has a Type {I, II, III} job, we express this customer's utility for that job as a concave, quadratic, and twice continuously differentiable separable function, $U_j(\cdot)$, defined as follows.

$$U_j(vm_j^{type}) = \alpha_j^{type} vm_j^{type} + \frac{\beta_j^{type}}{2}(vm_j^{type})^2, \tag{5}$$

where $type \in \{I, II, III\}$. vm_j^{type} is the amount of VM instances required to process j's entire job for a given type. Similar to the motivation and rationale behind the concave, quadratic cost functions for SCs, the utility function of an SC customer is designed such that the marginal utility for the customer is a *decreasing linear function* of the number of VM instances, i.e., the additional utility increase due to a unit increase in the number of VMs varies in a negative linear fashion with the number of VMs. α_j^e (a positive quantity) and β_j^e (a negative quantity) in the above equation are j's utility coefficients. *The above quadratic form of the cost function, apart from satisfying the property of negative linear marginals, not only allows for tractable analysis, but also serves as a good second-order approximation for the broader class of concave payoffs [17].*

With regard to modeling the VM instances for the different job types, we have (i) vm_j^e - the amount of VM instances required to process j's jobs of Type I, i.e., entire job, (ii) vm_j^c - the amount of VM instances required to process j's jobs of Type II, i.e., curtailed job, and (iii) vm_j^s - the amount of VM instances required to process j's jobs of Type III, i.e., time-shiftable tasks of a Type III job. Note that the non time-shiftable tasks of a Type III job can be considered as a job of Type I. vm_j^c again is expressed as an additive function of vm_j^e in the following manner:

$$vm_j^c = \kappa_j^1 vm_j^e + \kappa_j^2 vm_j^e,\ \kappa_j^1, \kappa_j^2 \in (0, 1).$$

Here, α_j^e (a positive value) and β_j^e (a negative value) are j's utility coefficients for Type I jobs. The interpretation of vm_j^c is as follows: $\kappa_j^1 vm_j^e$ is the number of VMs required to accomplish j's curtailed task, whereas $\kappa_j^2 vm_j^e$ is the additional number of unused VMs that contribute to j extra utility when its job is curtailed and provides it with an overall *perceived* satisfaction greater than that obtained from the utility derived solely using $\kappa_j^1 vm_j^e$ used VMs for the curtailed job.

A customer j can have jobs of all three types. Thus, his aggregate tasks are worth $vm_j^{ag} = vm_j^e + vm_j^c + vm_j^s$ VM instances. Therefore, customer j's aggregate utility takes a similar form to his utility for a specific job type, and is given by

$$U_j(vm_j^{ag}) = \alpha_j^{ag} vm_j^{ag} + \frac{\beta_j^{ag}}{2}(vm_j^{ag})^2, \tag{6}$$

where α_j^{ag} (a positive quantity) and β_j^{ag} (a negative quantity) are j's utility coefficients for his job aggregate. In the remainder of the paper, we will work with the aggregate number of VMs for a customer.

We denote π_j^{ag} to be the net utility that customer j generates through getting service for his aggregate job type from its contracted SC, and define the maximum net utility that customer j can generate, via the following optimization problem.

$$\max_{vm_j^{ag}} \pi_j^{ag} = \max_{vm_j^{ag}}[U_j(vm_j^{ag}) - \rho_j vm_j^{ag}]$$

subject to

$$vm_{\min_j^{ag}} \leq vm_j^{ag} \leq vm_{\max_j^{ag}}.$$

Here, $vm_{\min_j^{ag}}$ and $vm_{\max_j^{ag}}$ are the lower and upper bounds for the number of VM instances used up by customer j's job aggregate job type. ρ_j is the price paid by customer j to his chosen SC per VM instance used for his aggregate job.

2.3 The Regulator and Modeling Social Welfare

The Regulator's Role - The role of the regulator (e.g., the government, a federated agency) as applicable to our work is to ensure (i) good privacy practices between SCs, (ii) the design of policies/mechanisms that enable autonomous SCs to price customers appropriately without making excessive profits through market exploitation, and (iii) a appropriate level (potentially optimal) of social welfare allocation amongst the autonomous SCs at market equilibrium. (i) is specific to our problem setting and is the most important motivation for the presence of a regulator (see Section 1) in the first place. [1] However, the presence of a regulator brings in other important benefits through (ii) and (iii). (ii) is necessary to prevent any SC from exploiting its customers on service costs. In this work we *do not* focus on the design of such mechanisms, and assume the existence of one[2], whereas (iii) is important from an economic perspective as achieving *maximum* social welfare is a standard benchmark objective in welfare economics because it leads to (a) a certain level of equitability of allocations (in resources or in net utility) amongst the stakeholders, and (b) might guarantee *Pareto efficiency* at market equilibrium [12], and (c) an optimal social welfare state denotes the best possible operating point of an economic system. A *Pareto efficient* allocation of utilties amongst a set of stakeholders ensures that at market equilibrium none of the stakeholders can increase their net utility without decreasing any other stakeholder's net utility. The notion of equitability is important in the context of autonomous SC markets because they often operate in a decentralized fashion, and ideally, we would want a social welfare allocation at market equilibrium that does not result in considerable disparity amongst the players' allocations, something that can be ideally enforced in a centralized setting. In case of considerable disparity between the optimal social welfare and the welfare achieved in the decentralized market setting, the regulator can step in with policies that narrow the disparity gap.

Modeling Social Welfare - In this paper, we define the social welfare function of the regulator to be the sum of the net utilities of the SCs and their customers. We denote this function by SW, and express it as

$$SW = \sum_{j \in C} U_j(vm_j^{ag}) - \sum_{i \in SC} \left(c(vm_i^r) + c(vm_i^b) + c(vm_i^{pc}) \right), \tag{7}$$

where C is the set of consumers, SC is the set of small data centers, the first term is the sum of the utility of the consumers, and the second term is the sum of the costs faced by the SCs in SC for servicing customer demands. The aforementioned social welfare expression is the standard Bergson-Samuelson *utilitarian social welfare function* in social welfare economics [12] whose optimality does not necessarily focus on equality of resource or utility allocations amongst each class of stakeholders, i.e., the SCs and the customers, but only on Pareto efficiency of resource allocations amongst the stakeholders. *Note that in our autonomous SC setting, the regulator can only expect to have the social welfare function maximized in the best case because it cannot directly enforce optimal strategy choices of the SCs like in a centralized control setting*. The important question here is whether the utilitarian social welfare function is most relevant to this work.

We choose to work with the utilitarian function over two other popular Bergson-Samuelson social welfare functions used in economic applications: the *egalitarian function*, and the *Rawl's function*. The problem of optimizing the egalitarian function focusses on maximizing the sum of net utilities of SCs and their customers constrained on the equitable distribution of the resources within each class of stakeholders, whereas the problem of optimizing the Rawl's function focusses on maximizing the minimum resource allocation within each class of stakeholders constrained on the equality of utilities amongst each class of stakeholders. Optimizing the egalitarian or the Rawl's function does not necessarily result in Pareto efficiency. In this regard, for modeling purposes, we feel neither of these two social welfare functions to be appropriate for our problem setting as ensuring equality of resources or utilities amongst stakeholder classes along with Pareto efficiency is too strict of a requirement and practically reduces the feasible allocation solution space significantly. On the other hand, just aiming for equality of resources or utility amongst stakeholder classes without Pareto efficiency is not very practical for a heterogenous SC setting such as ours because stakeholders in non-Pareto efficient equality states always have the incentive to improve their utilities without decreasing anyone else's (e.g., via cooperation, collusion). Thus, we choose the utilitarian social welfare function whose optimization necessarily leads to Pareto efficient outcomes [12]. However, in Section 5, we do compare our results obtained via the utilitarian social welfare model, with that obtained from the egalitarian and the Rawlsian social welfare models.

We emphasize here that a federated agency can impose centralized control on the SCs (thereby breaking the autonomy of the SCs) to always achieve a socially efficient Pareto optimal allocation amongst stakeholder classes. The challenge in an autonomous SC setting is to achieve the centralized socially efficient Pareto optimal solution in a decentralized manner or getting as close as possible to such a solution. One of our primary goals in this paper is to address this challenge.

3 Static Market Analysis

In this section we derive and analyze perfectly competitive SC market equilibria. We assume that the autonomous SC firms decide on their (customer prices, quantity of reserved VMs by SCs) via a price-quantity competition game (e.g., a Cournot game [12]). We choose (price, quantity) as the parameters of competition among the SCs due to their publicly observable nature. Since prices in perfect competition are *strategic complements* (in the terminology of Bulow, Geanakoplos and Klemperer [19]), i.e., the decrease in an SC's customer price results in the decrease of customer prices charged by other SCs in competition, we are going to eventually converge to a stage where a single uniform customer price will prevail in the SC market [19]. We are interested to know whether such a price (also quantity) results in social welfare optimality. Equivalently, if a federated agency were to centrally impose a customer charging price (quantity) on all SCs (thereby breaking their autonomy) that would maximize social welfare, what would be the relationship between such a price (quantity) and the market equilibrium price (quantity) outcome of the price-quantity competition game? In this regard, we (a) formulate and solve an optimization problem for a regulator who wishes to achieve socially optimal market equilibria that maximizes utilitarian social welfare amongst the market stakeholders, (b) characterize market equilibria in the absence of a regulator and draw relationships between the equilibria obtained with socially optimal market equilibria. In practice, the competition between SC firms is likely to be imperfect in nature, and Laffont and Tirole have addressed models [20] under such settings which result in market efficiency. We consider a perfect competition setting in our work for the purpose of simplifying market equilibrium stability analysis in Section 4.

1. In practice, using mechanism design theory, the regulator can devise efficient economic mechanisms that enable SCs to find it incentive compatible in protecting the privacy of their customers. However, we do not focus on the design of such mechanisms in this paper, and subsequently do not address privacy modeling aspects.

2. Economists Laffont and Tirole have proposed *principal-agent* models in this regard [18] which will enable autonomous SCs to charge appropriate prices to customers purely out of self-interest.

Optimization Problem Formulation Here, we formulate a regulator's optimization problem so as to achieve socially optimal market equilibria. The primary goal of the formulation is to maximize the net utilities for the SC customers, and minimize the net cost of operation of SCs to reach a net maximum social welfare situation amongst the SCs and their customers. We define this problem mathematically as follows:

$$\textbf{OPT:} \quad \max SW$$

subject to

$$\sum_{j \in C_i} vm_j^{ag} - \sum_{i \in SC} (vm_i^r + vm_i^b + vm_i^{pc}) = 0, \, \forall i \in SC,$$

where the objective function is to maximize social welfare SW (see Equation 9 above) or equivalently to minimize the negative of social welfare (to have a convex objective function to fit the convex programming paradigm), and the constraint is the supply-demand balance equation, with $\sum_{j \in C_i} vm_j^{ag}$ representing total customer demand, and $\sum_{i \in SC}(vm_i^r + vm_i^b + vm_i^{pc})$ representing total SC supply. C_i is the set of customers served by SC i. A potential solution to the above optimization problem indicates the parameters at which the SC market can ideally operate and (i) make all stakeholders satisfied to a point that no one has an incentive to deviate, and (ii) maximize the total satisfaction of all the stakeholders together. We denote such an ideal state of market operation as a *static socially efficient market equilibrium.*

Dual Problem Formulation We will solve OPT using the *primal-dual* approach [21]. The advantage of using the primal-dual approach is that the dual optimization problem of the primal is always convex [21], and its solution results in global optima which can be related back to the optimal solution of the primal problem. Before deriving the dual optimization problem, we first define the Lagrangian function of OPT as follows:

$$\begin{aligned} L &= \sum_{i \in SC} \Big(c(vm_i^r) + c(vm_i^b) + c(vm_i^{pc}) \Big) - \sum_{j \in C} U_j(vm_j^{ag}) \\ &+ \left(\sum_{i \in SC} \rho_i \sum_{j \in C_i} vm_j^{ag} - \sum_{i \in SC} \rho_i (vm_i^r + vm_i^b + vm_i^{pc}) \right), \end{aligned}$$

where $\rho = (\rho_1, ..., \rho_n)$ is the vector of Lagrange multipliers for the constraint in OPT. The dual optimization problem, DOPT, is then defined as follows.

$$\textbf{DOPT:} \quad \max \inf_{\{vm^e, vm^c, vm^s, vm^r, vm^b, vm^{pc}, \rho\}} L,$$

where vm^e, vm^c, and vm^s are vectors of customer VM types and vm^r, vm^b, and vm^{pc} are vectors of SC VM types. Note that vm_i^{ag} for any customer i equals $vm_i^e + vm_i^c + vm_i^s$. Thus, the goal here is to find an optimal $\{vm^e, vm^c, vm^s, vm^r, vm^b, vm^{pc}, \rho\}$ tuple that is an optimal solution to both OPT and its dual.

Solving the Dual The dual optimization problem is convex and its optimal solution is found by applying the *Karush-Kuhn-Tucker* (KKT) conditions [21] that are stated through equations (10a)-(10e). Solving these equations, we obtain the optimal solution to DOPT. Since OPT is convex, applying *Slater's conditions* we obtain strong duality, i.e., a duality gap of zero[21], which implies that the optimal solution to OPT coincides with that of DOPT, and there is no loss in the value of the optimal solution by the transformation of the primal problem to its dual. The optimal solution to OPT/DOPT is unique, and is the static market equilibrium. We denote this solution by the tuple $\{vm^{e*}, vm^{c*}, vm^{s*}, vm^{r*}, vm^{b*}, vm^{pc*}, \rho^*\}$. We now state the KKT conditions in the form of equations 8(a-g) as follows.

$$\frac{d(c(vm_i^r))}{dvm_i^r} | vm_i^{r*} - \rho_i^* = 0, \, \forall i \in SC. \tag{8a}$$

$$\frac{d(c(vm_i^b))}{dvm_i^b} | vm_i^{b*} - \rho_i^* = 0, \, \forall i \in SC. \tag{8b}$$

$$\frac{d(c(vm_i^{pc}))}{dvm_i^{pc}} | vm_i^{pc*} - \rho_i^* = 0, \, \forall i \in SC. \tag{8c}$$

$$\rho_i^* - \frac{\partial(U_i(vm_i^e))}{\partial vm_i^e} | vm_i^{e*} = 0, \, \forall i \in C. \tag{8d}$$

$$\rho_i^* - \frac{\partial(U_i(vm_i^c))}{\partial vm_i^c} | vm_i^{c*} = 0, \, \forall i \in C. \tag{8e}$$

$$\rho_i^* - \frac{\partial(U_i(vm_i^s))}{\partial vm_i^s} | vm_i^{s*} = 0, \, \forall i \in C. \tag{8f}$$

$$\sum_{j \in C_i} vm_j^{ag}(1 - \kappa_j^1 - \kappa_j^2) = (vm_i^r + vm_i^b + vm_i^{pc}), \, \forall i \in SDC. \tag{8g}$$

Thus, we have our first result - *SC markets have a unique static socially efficient market equilibrium.* The uniqueness of the result is due to the convexity of the dual formulation. This unique equilibria can always be achieved by the regulator in a centralized manner by breaking the autonomy of the SCs. *The key question is whether such an equilibria can be realized in an autonomous setting.* Based on the general equilibrium theory in microeconomics [12], market equilibria in a perfectly competitive setting of firms is known as *Walrasian equilibria*. It turns out from general equilibrium results in [12] that the unique optimal solution to OPT (i) is a competitive Walrasian equilibrium that is Pareto efficient, (ii) satisfies *Arrow-Debreu's* first and second fundamental theorems of welfare economics that establishes the *if and only if* relation between the existence of a Walrasian equilibrium and its Pareto efficiency [12], (iii) maximizes utilitarian social welfare, and (iv) clears the market by balancing total SC resource supply with consumer and SC resource demand. *We consider this unique equilibrium state to be the benchmark at which the SC market would be willing to always operate.*

In practice, for a perfectly competitive market, there may be multiple Pareto efficient Walrasian market equilbria that are not socially efficient. In Section 5, we compare the social welfare performance of such equilibria induced by the market setting with the optimal social welfare derived from the socially efficient market equilibrium.

Computing the Socially Optimal Market Equilibrium The optimal solution to the dual optimization problem, DOPT, can be obtained in an iterative manner using a gradient approach, the principle behind which is the *Primal-Dual Interior Point Method* [21]. We adopt the Primal-Dual Interior Point method in our work because it has a polynomial-time complexity to arrive at the optimal solution to convex programs [22]. The basis of the method is to *progressively* change the argument vector of DOPT so that minima-Lagrange multiplier ρ satisifes the KKT conditions.

Denote by v, DOPT's argument vector sans the Lagrange multiplier ρ, $\{vm^e, vm^c, vm^s, vm^r, vm^b, vm^{pc}\}$. Applying the Interior Point method to DOPT gives us the the following equations:

$$v(t + \epsilon) = v(t) - k_v \nabla_x L \cdot \epsilon. \tag{9a}$$

$$\rho(t + \epsilon) = \rho(t) + k_\rho \nabla_x L \cdot \epsilon. \tag{9b}$$

Here, k_v and k_ρ are positive scaling parameters which control the amount of change in the direction of the gradient. Letting

$\epsilon \to 0$, we get

$$\tau_v \dot{v}(t) = -\nabla_v L, \tag{10a}$$
$$\tau_v \dot{\rho}(t) = -\nabla_\rho L, \tag{10b}$$

where $\tau_y = \frac{1}{k_y}$ for $y = v, \rho$. The Interior Point Method converges in polynomial time when the duality gap approaches zero, due to the linear and super-linear convergence rate of the method [21].

4 Dynamic SDC Markets

In the previous section, we analyzed static SC markets. We showed that autonomous SCs in a perfectly competitive setting are able to eventually converge to the static market equilibrium point which is also the socially optimal welfare state. In this section, we study the stability of an SC market to remain in the socially optimal (also market efficient) state. In practice, an SC market can be dynamic in nature due to the non-static nature of the supply of data center resources and variability of over time of customer demand. This dynamic nature of the SC market is likely to lead to frequent static market equilibrium perturbations, which in turn might (not always) lead to a state of market disequilibrium. Here, the term 'disequilibrium' refers to a state when market supply does not equal market demand due to perturbations in market parameters (e.g., customer prices), and as a result all stakeholders do not mutually satisfy their interests. *In such a case, an important challenge is to design a stable market that is robust to perturbations and always returns to its equilibrium point(s) when market disequilibrium results.* Inspired by the notion of *disequilibrium process* [2], we propose a *dynamic market mechanism* for SCs. The concept of disequilibrium pertains to a situation where a static market equilibrium is perturbed, potentially to a disequilibrium state, and the underlying players (stakeholders) work together to re-attain the equilibrium. *The main idea behind the disequilibrium process is an iterative sequence of* action *and* state *profiles (see below), i.e., information exchange between the dominant market stakeholders, of VM instance supply and demand levels, and per-unit VM instance prices, to arrive at a desired static equilibrium.* Such an iterative process essentially implies an overall dynamic model with feedback. Our proposed dynamic market mechanism can also be used to re-attain a specific preferred equilibrium point from a given equilibrium point. We first present our dynamic market model and then follow it up with its stability analysis.

4.1 Dynamic Model

Our dynamic model of SC markets consist of a *state space*, $X \subset \mathbb{R}^n$, where each state, $\{\rho_i\} \in X$, is the profile of per-unit VM instance prices at each SC i. The *state dependent payoff*, i.e., profit function for each SDC from its reserved resources is given by

$$\pi_i^r = \rho_i vm_i^r - c(vm_i^r).$$

The state dependent payoff for each SC from its borrowed resources is given by

$$\pi_i^b = \rho_i vm_i^b - c(vm_i^b).$$

Similarly, state dependent payoff for each SC from resources borrowed from a public cloud is given by

$$\pi_i^{pc} = \rho_i vm_i^{pc} - c(vm_i^{pc}).$$

The payoff function for the SC customers for a given job type $\in \{e, c, s, ag\}$, is given by

$$U_j(vm_j^{type}) - \rho_j vm_j^{type}.$$

Each SC is assigned a *state dependent action* that permits the SCs and their customers to change their VM instance generation and consumption levels respectively. We assume a *Cournot competition* [14] of VM instance prices amongst the SCs in competition, and following that the action for each SC i consists of commiting a certain amount of VM instances that influences the market-clearing process. In this paper, we use the *gradient play* in game theory [13] to derive the state dependent actions of the SCs and their customers. Our use of the gradient play technique is motivated by the fact that in many practical market environments stakeholders (i) find it behaviorally difficult or computationally expensive to play their *best responses* [14], (ii) have zero or incomplete knowledge of the utilities of other stakeholders in the market, and (iii) cannot even observe the actions of other stakeholders in the worst case. In such environments, gradient play is a suitable technique to achieve static market equilibrium stability iteratively [15]. More specifically, for our market setting the occurrence of (i)-(iii) is quite likely. Gradient play also works when issues (i)-(iii) do not arise. *The main idea behind the gradient play technique is the use of ordinary differential equations (ODEs) to describe the path of a perturbed system state to the static market equilibrium state.* Using gradient play, the action for the the ith SC is given by

$$\tau_i^r \dot{vm_i^r} = \rho_i - \beta_i^r vm_i^r - \alpha_i^r. \tag{11a}$$
$$\tau_i^b \dot{vm_i^b} = \rho_i - \beta_i^b vm_i^b - \alpha_i^b. \tag{11b}$$
$$\tau_i^{pc} \dot{vm_i^{pc}} = \rho_i - \beta_i^{pc} vm_i^{pc} - \alpha_i^{pc}. \tag{11c}$$

Here, the parameters τ_i^r, τ_i^r, and τ_i^r are time constants that describe the speed with which the action of VM instance commitment by SC i can be adjusted, and are free parameters to be determined. The goal of SC i's action is to drive the solution vm_i^r, vm_i^b, and vm_i^{pc} to vm_i^{r*}, vm_i^{b*}, and vm_i^{pc*}, the solution to Equations 8(a-c) at static market equilibrium. It can be seen that the RHSs of 11(a-c) are proportional to the gradient $\nabla_{vm_i^r} L$, $\nabla_{vm_i^b} L$, and $\nabla_{vm_i^{pc}} L$ respectively, where L is the Lagrangian of OPT. The suite of equations 13(a-c) can be solved independently by SC i. In a similar fashion, using gradient play, the state dependent action for any SC customer $i \in C$ is given by

$$\tau_i^e \dot{vm_i^e} = \beta_i^e vm_i^e + \alpha_i^e - \rho_i. \tag{12a}$$
$$\tau_i^c \dot{vm_i^c} = \beta_i^e vm_i^c + \alpha_i^c - \rho_i. \tag{12b}$$
$$\tau_i^{pc} \dot{vm_i^{pc}} = \rho_i - \beta_i^{pc} vm_i^{pc} - \alpha_i^{pc}. \tag{12c}$$

$\{\tau_i^e, \tau_i^c, \tau_i^s\}$ are free parameters to be determined that denote the speed with which the consumption action of SC customer i can be adjusted. The goal of the SC customer action here is to drive the solution $\{vm_i^e, vm_i^c, vm_i^s\}$ to $\{vm_i^{e*}, vm_i^{c*}, vm_i^{s*}\}$, the solution to Equation 8(d-f) at static market equilibrium. It can be seen that the RHS of equation suite 12 is proportional to the gradient $\{\nabla_{vm_i^e} L, \nabla_{vm_i^c} L, \nabla_{vm_i^s} L\}$, $i \in C$, where L is the Lagrangian of OPT. Equation 12 can be solved independently by each SC customer i.

The dynamics of the pricing mechanism can be expressed via the following equation.

$$\tau_{\rho_i} \dot{\rho_i} = \sum_{j \in C_i} vm_j^{ag}(1 - \kappa_j^1 - \kappa_j^2) - (vm_i^r + vm_i^b + vm_i^{pc}), \tag{13}$$

where the goal is to drive the solution ρ_i, $\forall i \in SC$ to ρ_i^*, the solution of 8(g) at static market equilibrium. Here, τ_{ρ_i} is the free parameter denoting the speed with which ρ_i can be adjusted.

Equations 11-13 represent a dynamic model of the overall SC market. It resembles a repeated negotiation process where SC i responds with a commitment of vm_i^x, $x \in \{r, b, pc\}$ to suggested prices ρ_i received from the regulator; SC customer i responds with a consumption amount of vm_j^{type}, type $\in \{e, c, s, ag\}$, to the same prices. The regulator in turn adjusts its prices to these actions by the SCs and their customers and returns new prices, $\{\rho_i\}$, and the process continues till

convergence to the static market equilibrium derived in Section 3.

4.1.1 A Compact Representation

We need to compactly represent the above dynamic SC market model to pave the way for analytically analyzing the stability of such markets via the *Arrow-Hurwicz* criterion that is based on the theory of Lyapunov stability (see Section 4.2). Using Equations 11-13, our proposed dynamic market mechanism can be compactly represented in the matrix form via the following equation:

$$\begin{bmatrix}\dot{x}_1(t)\\ \dot{x}_2(t)\end{bmatrix} = \begin{bmatrix}A_1+\Delta A_1 & A_2\\ 0 & 0\end{bmatrix}\begin{bmatrix}x_1(t)\\ x_2(t)\end{bmatrix} + \begin{bmatrix}\bar{\alpha}\\ f_2(x_1x_2).\end{bmatrix} \quad (14)$$

Definiton of Equation Parameters. We now describe the parameters of Equation 14. We have

$$x_1(t) = [VM^r_{SC}\, VM^b_{SC}\, VM^{pc}_{SC}\, VM^e_C\, VM^c_C\, VM^s_C\, \Delta\, \rho]^T$$

that is a vector of dimension $(|SC|+|C|+2|SC|-1)\times 1$. Here, $|SC| = n$. We also have

$$x_2(t) = [0]_{n-1\times 1},$$

and

$$A_1 = \begin{bmatrix}-M_1 & 0 & 0 & M_2\\ 0 & M_3 & 0 & -M_4\\ 0 & 0 & 0 & -M_5\\ -M_6 & M_7 & M_8 & 0\end{bmatrix}$$

,

$$A_2 = [0\,0\; -M_9\, 0].$$

We define matrices M_1 to M_9 as follows: $M_1 = Diag(\frac{1}{\tau_i^{type}}\beta_i^{type})$, type$\in \{r,b,pc\}$. We assume that all for a given type, τ_i^{type}'s are equal for all $i \in SC$. $M_2 = Diag(\frac{1}{\tau_i^{type}}A^T_{SC})$, type$\in \{r,b,pc\}$, where $A_{SC} = Diag(1)$. $M_3 = Diag(\frac{1}{\tau_i^{type}}\beta_i^{type})$, type$\in \{e,c,s,ag\}$. $M_4 = Diag(\frac{1}{\tau_i^{type}}A^T_C)$, type$\in \{e,c,s,ag\}$, where $A_C = Diag(1)$. $M_5 = Diag(A'^T BA)$, where A' is an $(n)\times(n-1)$ matrix of 1's except for the 0 diagonal elements, B is an $n\times n$ matrix with all entries 1 except for entries of the form B_{ii} that take a value of zero, and A is an $n\times n-1$ matrix. $M_6 = Diag(\frac{1}{\tau_{\rho_i}^{type}}A_{SC})$, type$\in \{r,b,pc\}$. $M_7 = Diag(\frac{1}{\tau_{\rho_i}^{type}}A_C)$, type$\in \{e,c,s,ag\}$. $M_8 = Diag(\frac{1}{\tau_{\rho_i}^{type}}A^T BA')$, where A is an $(n-1)\times n$ matrix. $M_9 = [1]_{n\times n}$.

The expression $f_2(x_1,x_2)$ is a projection function onto the non-negative orthant, and is given by

$$f_2(x_1,x_2) = [cx_1 - VM^{\max}]^+_{x_2}, \quad (15)$$

where $c = BA'R$, R being a rotating matrix. of dimensionality $((|SC|-1)\times|SC|+|C|+2|SC|-1)\times 1$, and $VM^{\max}$ denotes a vector of maximum VM instances committed by each individual SC. The nth row of the projection $[cx_1 - VM^{\max}]^+_{x_2}$ is denoted as

$$\left[[cx_1 - VM^{\max}]^+_{x_2}\right]_n = \begin{cases}\max(0,[cx_1]_n - VM^{\max}_n, & \text{if } [x_2]_n = 0\\ [cx_1]_n - VM^{\max}_n, & \text{if} [x_2]_n > 0\end{cases} \quad (16)$$

ΔA_1 in Equation 14 represents the resource availability perturbations due to dynamics of the SC market caused by factors stated in Section 4.1. The value lies in a perturbation set E, where E is given by

$$E = \{\Delta_A = \Delta_{SC} - \Delta_C | \Delta_{SC} \in E_{SC};\, \Delta_C \in E_C.\} \quad (17)$$

Here,

$$\Delta_{SDC} = \begin{bmatrix}M_{10} & 0 & 0 & 0\\ 0 & 0 & 0 & 0\\ 0 & 0 & 0 & 0\\ M_{11} & 0 & 0 & 0\end{bmatrix},$$

where matrix M_{10} is given by $Diag\left(\frac{1}{\tau_i^{type}}\beta_i^{type}(\Delta_{SC})^2\right)$, type$\in \{r,b,pc\}$, and $\Delta_{SC} = Diag(\Delta_{SC}^{type})$. Matrix M_{11} is given by $Diag\left(\frac{1}{\tau_{\rho_i}^{type}}A^T_{SC}(I-\Delta_{SC}^{type})\right)$, and $A_{SC} = Diag(1)$. We also have E_{SC} expressed via the following:

$$E_{SC} = \{\Delta_{SC} | ||\Delta_{SC}|| = \sqrt{\lambda_{\max}(\Delta^T_{SC}\Delta_{SC})} \le \pi_{SC}\},$$

where π_{SC} is a finite constant. Similar to the expression for Δ_{SC}, we have

$$\Delta_C = \begin{bmatrix}0 & 0 & 0 & 0\\ 0 & 0 & 0 & 0\\ 0 & 0 & 0 & 0\\ 0 & M_{12} & 0 & 0\end{bmatrix},$$

where the matrix M_{12} is given by $Diag\left(\frac{1}{\tau_{\rho_i}^{type}}A^T_C(I-\kappa_j^1-\kappa_j^2)\right)$. We also have

$$E_C = \{\Delta_C | ||\Delta_C|| = \sqrt{\lambda_{\max}(\Delta^T_C\Delta_C)} \le \pi_C\},$$

where π_C is a finite constant. Finally, we express $\bar{b}$ as

$$\bar{b} = \left[Diag(\frac{1}{\tau_i^x}\alpha_i^{type}) + Diag(\frac{1}{\tau_i^x}\alpha_i^x)\Delta_{SDC}^{type}\; Diag(\frac{1}{\tau_i^y}\alpha_i^{type})\; 0\right]^T,$$

where $x \in \{r,b,pc\}$, and $y \in \{e,c,s,ag\}$. We assume that for given x, y, the values of α_i^x and α_i^y are equal for all i.

4.2 Stability Analysis of Dynamic Markets

In this section, we derive results regarding the stability of static market equilibria in a dynamic SC market setting. Specifically, (i) we derive the dynamic market equilibria obtained through gradient play mechanics and compare it with the socially efficient static market equilibria obtained in Section 3, and (ii) study the region of attraction around dynamic market equilibria to derive stability connotations.

We *first* consider stability aspects when κ_j^1, κ_j^2 equals zero, i.e., there are no curtailed jobs. In this case, the equilibria of the dynamic SC market described through Equations 11a - 11c (via the use of the gradient play technique), lies in the set

$$E = \{(x_1,x_2) | A_1x_1 + A_2x_2 + \bar{\alpha} = 0 \cap f_2(x_1,x_2) = 0\}.$$

Let (x_1^*, x_2^*) be an equilibrium point in set E. We then have the following theorem stating the relationship between (x_1^*, x_2^*) and the unique static SC market equilibrium obtained through Equations 10(a) - 10(e). The proof of the theorem is in the Appendix.

Theorem 1. *The equilibrium (x_1^*, x_2^*) is identical to the unique static market equilibrium obtained from the solution of OPT.*

Theorem Implications. The theorem suggests that in the absence of curtailed jobs, the equilibrium in a dynamic market setting is unique, and converges to the static market equilibrium in which the market existed initially before it was perturbed. Intuitively, when the SC market is perturbed from its equilibrium setting, a disequilibrium state might result, which will get resolved due to our proposed gradient-play based approach that rolls back the disequilibrium state to the original socially optimal static equilibrium state. In this paper, we are able to roll back to the original state in theory because of our assumptions regarding

the nature of utility functions. In practice, gradient play will guarantee a roll back of a disequilibrium market state to an equilibrium state not necessarily the original equilibrium state from which it was perturbed.

We now investigate the stability of the dynamic market equilibrium to find the region of attraction around itself. We introduce a few definitions in this regard. Let $y_1 = x_1 - x_1^*$, $y_2 = x_2 - x_2^*$. Denote by $V(y_1, y_2)$ a scalar, positive definite Lyapunov function expressed as

$$V(y_1, y_2) = y_1^T y_1 + y_2^T P_2 y_2, \tag{18}$$

where P_1 and P_2 are diagonal matrices. We use *Lyapunov functions* from control theory [23] as a standard to prove the stability of an equilibrium of a system represented via ordinary differential equations (ODEs), such as the ones arising in our work in Section 4.1. Let d be expressed as

$$d = \frac{2\lambda_{min}(P_2)\psi_{\min}\lambda_{\min}(Q)}{\beta^2}, \tag{19}$$

where $\lambda_{\min}(\cdot)$ denotes the minimum eigenvalue of Q,

$$\beta \geq ||P_1 A_2 + R^T[1]_{n\times n} P_2||_2,$$

where R is a rotating matrix, and $\psi_{\min} = \min(\psi_i)$, ψ_i being the coefficient of the orthogonal vector w_i to express $VM^{\max}$ as $\sum_{i=1}^{n} \psi_i w_i$. We now have the following theorem characterizing stability of the dynamic market equilibrium. The proof of the theorem is in the Appendix.

Theorem 2. *Let A_1 be Hurwitz. Then the equilibrium (x_1^*, x_2^*) is asymptotically stable for all initial conditions in*

$$\Omega_{c_{max}} = \{(y_1, y_2) || V(y_1, y_2) \leq c_{max}\} \text{ for } c_{max} > 0,$$

such that

$$\Omega_{c_{max}} \subsetneq D = \{y_2 \geq 0 | ||y_2||_2 \leq d\}$$

.

Theorem Implications. Intuitively, the theorem states that irrespective of any initial state the market is in, on being perturbed, it will always come back to an equilibrium state from a disequilibrium state. The *Hurwitz* (not the same as Hurwicz) nature of matrix A_1 is determined from the time constants in Equations 11-13. Most real systems satisfy the Hurwitz criterion in that A_1 will be a *real square matrix* constructed with coefficients of a real polynomial.

We *now* consider stability aspects when κ_j^1, κ_j^2 *does not* equal zero. In this case, the equilibria of the dynamic SC market described through Equations 11a - 11c, also lies in the set E. We define y_1, y_2, and $V(y_1, y_2)$ as before but define d_Δ as

$$d_\Delta = d - d_{\Delta_{SC}} + d_{\Delta_C}, \tag{20}$$

where d is the same as in Equation (21), Δ_{SC} and Δ_{SC} represent the supply demand perturbation matrices, and $d_{\Delta_{SC}}$ and d_{Δ_C} are given by

$$d_{\Delta_{SC}} = \frac{4\lambda_{min}(P_2)\psi_{\min}||P_1||_2 \pi_i | i \in SC}{\beta^2}. \tag{21a}$$

$$d_{\Delta_C} = \frac{4\lambda_{min}(P_2)\psi_{\min}||P_1||_2 \pi_j | j \in C}{\beta^2}. \tag{21b}$$

We now have the following theorem characterizing market stability. The proof of the theorem is in the Appendix.

Theorem 3. *Let A_1 be Hurwitz, and let*

$$\pi_{SC} - \pi_C < \frac{\lambda_{\min}(Q)}{2||P_1||_2} \tag{22}$$

Then the equilibrium (x_1^*, x_2^*) is asymptotically stable for all initial conditions in

$$\Omega_{c_{max}} = \{(y_1, y_2) || V(y_1, y_2) \leq c_{max}\} \; for \; c_{max} > 0,$$

such that $\Omega_{c_{max}} \subsetneq D = \{y_2 \geq 0 | ||y_2||_2 \leq d_\Delta\}$.

Theorem Implications. Similar to the implications of Theorem 2, this theorem states that irrespective of any initial state the market is in, on being perturbed, it will always come back to an equilibrium state from a disequilibrium state.

5 Numerical Evaluation

In this section, we numerically evaluate our dynamic market model to investigate (i) market equilibrium (ME) performance of static markets with different social welfare (SW) metrics and (ii) stability behavior of dynamic markets. The first part of this section describes the evaluation setting, and the second part analyzes the evaluation results.

5.1 Experimental Setup

As a representative experimental setting, we consider five SCs and 15 customers (not including other SCs). Each SC has five customers each and they are tied to the SCs throughout the entire duration of the experiment. Peer SCs are assumed to be altruistic w.r.t. VM borrowing. The market parameters for the SCs and the customers are shown in Tables 1 and 2 respectively. We simulate a Cournot price-quantity competition game between the SCs, and use the *tatonnement process* (TP) [24] to converge to a static market equilibrium in practice for a distributed setting. *Tatonnement* is a trial-and-error process similar to the *hill climbing* approach in local search theory by which equilibrium is reached in competitive markets via a distributed fashion. As a measure of static market efficiency we investigate and compare the SW function values at ME for *utilitarian, egalitarian*, and *Rawlsian (see Section 2.3 for more details)* markets with optimal SW values. For the parameter values in Tables 1 and 2, we run experiments for all possible permutations (instances) of values that are applicable to SCs and their customers, and report the mean value of the results obtained *(with the exception of Figure 2c which reports on individual permutations)*. Note that each permutation of values can be considered as a different market setting. To experiment on dynamic markets, we fix τ_ρ to be the same for all SCs and vary it in the interval $[0, 5]$. Similarly we fix τ^{ag} to be the same for all 15 customers and vary it in the interval $[.05, .2]$. We also make κ^1 and κ^2 to be equal for all customers and vary it in the interval $[0, 0.05]$. For the purpose of simulation, we relax the assumption in Section 2 on the quadratic concave cost functions for SCs and SC customers, and also experiment on non-quadratic concave cost functions with negative marginals. More specifically, with respect to Equation 1, we consider $f_1(\cdot)$ to be a linear function, and $f_2(\cdot)$ to be a non-linear non-quadratic function with degree strictly greater than 2; in our experiments we take degree values of 2.5, 3, and 4.

TABLE 1: Parameters of Cost Functions for SCs

SC#	$VM^{\min}$	$VM^{\max}$	τ_{SC}	β	α
SC1	0	200	0.6	-0.3	90
SC2	0	200	0.2	-0.6	102
SC3	0	250	0.6	-0.25	80
SC4	0	250	0.6	-0.25	80
SC5	0	200	0.2	- 0.01	20

TABLE 2: Parameters of Utility Functions for Customers

Customer(C)#	$VM^{\min}$	$VM^{\max}$	τ_C	β	α
C1, C2	60	100	0.1	-0.5	168
C3, C4	60	100	0.1	-0.15	140
C5, C6	70	80	0.2	-0.35	140
C7, C8	20	60	0.2	-0.2	100
C9, C10	30	60	0.2	-0.3	120
C11, C12	20	40	0.2	-0.1	125
C13, C14, C15	30	60	0.2	-0.5	135

5.2 Analysis of Experimental Results

In the first part of this section, we analyze SC cost and customer utility allocations at market equilibrium under the utilitarian, egalitarian, and Rawlsian SW paradigms. In the second part, we analyze the stability of various dynamic market settings, and also how fast a dynamic market converges to a stable equilibrium.

Static Market Equilibrium Performance. Using TP, for each SW paradigm, we arrive at a different single market equilibrium maximizing SW. We note here that the MEs obtained via the tatonnement process are not necessarily the best, if the market has multiple equilibria. *We observe from Figures 1, 2 (a&b) that with respect to SC and customer allocation ratio fairness, Rawlsian MEs are the best as they ensure nearly identical cost and utility allocations across all autonomous SCs and customers respectively, followed closely by egalitarian MEs, and utilitarian MEs that are not very fair.* Here, we define *allocation ratio* as the ratio of the cost (utility) of SC (customer) i to the minimum (maximum) cost (utility) of any SC (customer) at ME. *On the other hand, we see that market equilibrium in utilitarian markets, MEs lead to a considerably greater maximum social welfare (see Figures 1c. and 2c.) when compared to egalitarian and Rawlsian markets.* In addition, from theory, SW maximizing ME in utilitarian competitive markets are *always* Pareto optimal. In Figures 1c and 2c, $SWOPT_t$, $t \in \{U, E, R\}$, denotes the maximum SW under a given social welfare paradigm t, and SW_t^{ME} denotes the SW value at ME of a given social welfare paradigm t.

We also study the normalized fraction of resources, i.e.,VMs, in the *reserved*, *borrowed*, and *public cloud* categories for each SC with respect to its sum total of VMs required, for both quadratic and non-quadratic cost functions, at market equilibrium. Table 3 reports the case for quadratic SC cost functions, and Table 4 reports the case for non-quadratic cost functions. As a representative example, we choose to report results when SC cost functions have degree 4. The results for the cases when the cost functions have degrees 2.5 and 3 show very similar trends. *At market equilibrium, we observe that the normalized fraction of resources requested from the public cloud is the least for all SCs when compared to their fractions of reserved and borrowed resources.* This validates the effectiveness of our model to ensure that minimum requests are sent to the public cloud, a primary aim of our research. *We also observe that for each SC, the number of borrowed VMs are less than their reserved number of VMs.* This result implies that none of the SCs take undue advantage of peer SCs in borrowing more than they can reserve, and also do not find it cost effective to do so.

TABLE 3: Normalized Fractions of Resources for SCs (by category) at ME Quadratic Cost

SC#	reserved	borrowed	public cloud
SC1	0.55	0.35	0.1
SC2	0.60	0.25	0.15
SC3	0.68	0.20	0.12
SC4	0.68	0.20	0.12
SC5	0.77	0.15	0.08

TABLE 4: Normalized Fractions of Resources for SCs (by category) at ME Non-Quadratic Cost

SC#	reserved	borrowed	public cloud
SC1	0.50	0.38	0.12
SC2	0.57	0.25	0.18
SC3	0.62	0.19	0.19
SC4	0.62	0.19	0.19
SC5	0.72	0.16	0.12

Dynamic Market Stability Performance. Through Figures 3(a-c), we study dynamic markets for three different instances of (κ_1, κ_2) pairs, with quadratic cost functions. *We note here that we studied dynamic markets for the same three instances of (κ_1, κ_2) pairs, with non-quadratic cost functions having degrees 2.5, 3, and 4 respectively, and the results are very similar to the case with quadratic cost functions. Thus we only report plots for the quadratic case, as a representative example.* For each instance, *we observe that low values of τ_ρ for a given instance correspond to market instability, i.e., a state of disequilibrium, because they imply a fast update in SC prices charged to customers, indicating market volatility in supply and demand as well.* Here, stability is indicated through the maximum of the eigenvalues of Hurwitz matrix A_1 (see Section 4.1.1) formed from the market instance, which are negative in the stable zone, and positive in the unstable zone. It is logical to expect that market instability can be reduced if the price update is slower, i.e., if τ_ρ is larger. *We also observe that market volatility is increased due to a decrease in τ^{ag} values because the latter trend corresponds to the increase in demand elasticity which contributes to a market being volatile.* We infer from the plots that it is possible to design an SC market where volatility (arising due to either low τ_ρ or low τ^{ag} values) can be contained by increasing market latency, i.e., increasing τ_ρ values. With respect to the speed of convergence, from Figure 3, *we observe in general that SC markets converge fast to the stable zone, i,e., even at low values of τ_{rho}, but the speed of convergence increases with increasing κ_1, κ_2 values.* This is because increasing κ values indicate more demand curtailment by SC customers, potentially leading to non-volatility in supply and prices.

6 Related Work

Related work specific to SCs is quite recent and thus relatively scarce. In this section, we study works related with the SC setting and as applicable to our paper. In this regard, we divide this section into two parts: (i) works based on SC environments, and (ii) regulatory and non-regulatory control in cloud sharing.

6.1 Works Targeting SC Environments

In [25, 26], the authors have focused on effective platform design for enabling cooperation between cloud providers. Other works in the SC space have focussed on prices and revenue maximization for incentivizing data centers to cooperate [9]. However, in a dynamic setting of data centers with varying workload patterns and changing sociopolitical and legal realities, prices and revenue maximization are not the sole (or at times even relevant) criterion that can determine cooperation. Factors such as location proximity, workload similarity, privacy concerns etc., can be equally, if not more, important considerations for resource exchange and allocation between small data centers. In this regard, the authors in [1] propose a decentralized cloud with a swarm of networked SCs which cooperate with each other under varying workloads, and which employ various strategies in order to cooperate with others. Their analysis aims to answer the following questions: i) what are the incentives of SCs for cooperation within the decentralized cloud?, ii) how is the performance of cooperation affected

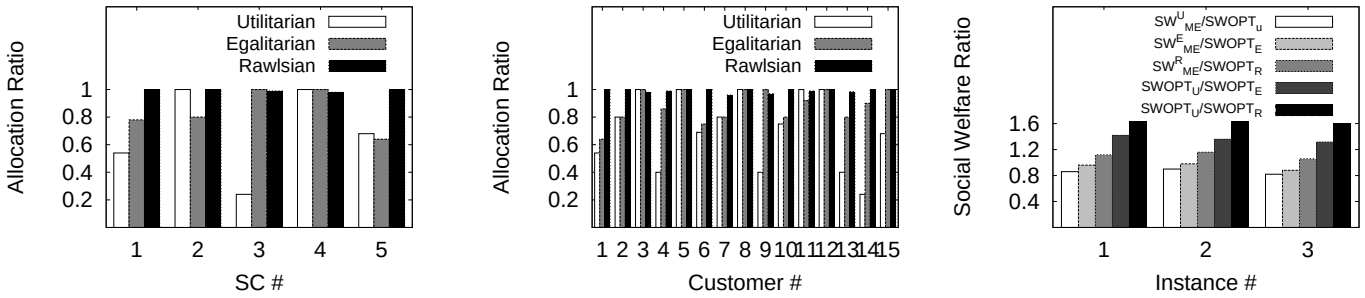

Fig. 1: ME Performance (a) SC Costs (left), (b) Customer Utilities (middle), (c) Social Welfare Ratio (right) w.r.t to various SW metrics (Quadratic costs)

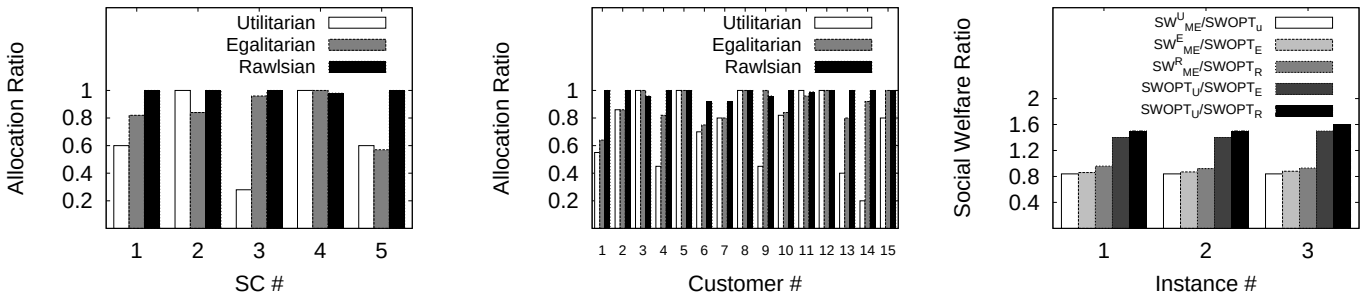

Fig. 2: ME Performance (a) SC Costs (left), (b) Customer Utilities (middle), (c) Social Welfare Ratio (right) w.r.t to SW metrics (Non-Quadratic Costs)

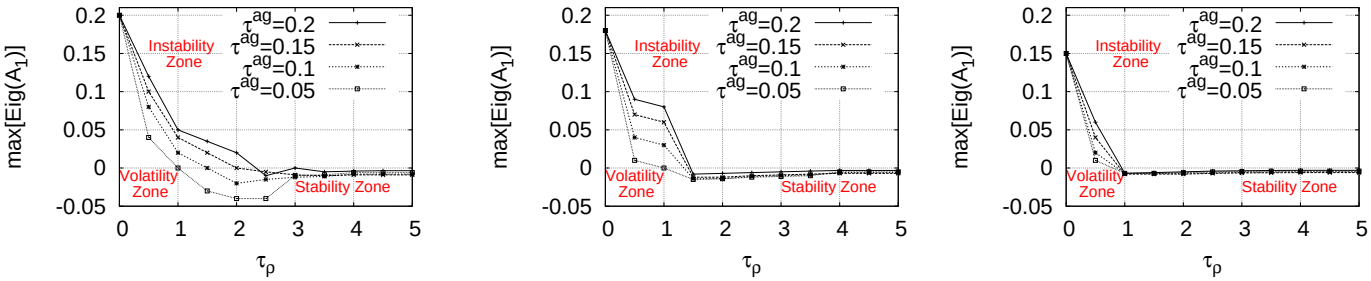

Fig. 3: Market Stability Performance when (a) $\kappa_1 = \kappa_2 = 0$ (left), (b) $\kappa_1 = \kappa_2 = 0.02$ (middle), (c) $\kappa_1 = \kappa_2 = 0.05$ (right)

due to different strategies adopted by SCs?, and (iii) what type of strategy can thrive in a heterogeneous environment?. Subsequently, the authors design general strategy functions (based on different dimensions: capacity, history, prediction and reciprocity, that can be employed by SCs to evaluate the performance of cooperation between the SCs. Their results reveal that most of the SCs will eventually converge on the same strategy in order to achieve the state of stable mutual cooperation.

Differences with Our Work. Our work is both an extension as well as orthogonal to the papers mentioned above. It is an extension in the sense that we model costs for borrowing SC resources from peers, whereas the above mentioned works assume free resource sharing between SCs. It is orthogonal to the above mentioned approach in that we study robust supply-demand markets for SCs, given an effective platform for enabling cooperation between cloud providers. In that regard, we investigate the optimal parameters (see Section 2) at which SC markets should operate so as to jointly satisfy market stakeholders, i.e., the supply and the demand sides, in an economically efficient manner.

Relationship of Our Work with Network Markets. A recent work[27] in existing literature is the first to address markets specifically pertaining to network applications. In [27], the authors investigate a service exchange model for exchange networks and analyze the network users' interactions under three different approaches: (i) a central allocation policy decided by federated agency, (ii) a competitive market approach, and (iii) a coalitional game between the network users. The model considered in [27] is generic and representative of many communication or economic networks. As the main contribution, the authors in [27] fully characterize the market equilibria for cases (i)-(iii), relate them to the *lexicographically max-min* fair solution, and study how the equilibria are affected by the network graph. *However, unlike us, the model in [27] does not (i) presume any type of money transfers, i.e., there is no budget constraints (as in typical exchange economies) and the market players do not value money, (ii) do not target reaching the social welfare optimal solution, which is atleast as good as the max-min fair solution and in most cases better, and (iii) do not address market instability issues that might arise to resource fluctuations.*

6.2 Control in Data Center Sharing

Aspects of regulatory (federated) and non-regulatory control of small clouds are important in the context of effective resource sharing management among clouds.

Non-Regulated and Decentralized Control In [25], the authors propose a non-regulated decentralized cloud resource sharing mechanism that uses volunteer edge resources for data intensive computing applications. Wang et al. [28] study a decentralized non-regulated cloud platform, *SpotCloud*, built on customer-provided resources that allows customers to sell their idle resources to offer cloud services collaboratively. Another approach proposed by the authors in [26] designs and implements a decentralized non-federated cloud resource sharing based computing platform at an Infrastructure as-a-Service (IaaS) level, based on P2P technologies. In [1], the authors focus on studying the incentives and strategies of resource sharing among non-federated SCs and on how the SCs can cooperate in a decentralized manner.

Regulated and Centralized Control Early approaches to a model of a cloud federation can be found in [29, 30, 31]. Buyya et al. [32] introduced a market oriented VM exchange model among clouds. Federations of hybrid clouds were also discussed in [33, 34]. Goiri et al. [33] proposed a novel federation management architecture and focused on building an accurate revenue function of CPs when taking one of the following decisions: (i) participate in the federation by outsourcing or insourcing capacity or (ii) turning off spare capacity. Similarly, Toosi et al. [34] provided a comprehensive analysis of the related costs and revenue associated with the decisions in (i) and (ii), of CPs in the federation. A mechanism to dynamically allocate resources of distributed data centers in a federated centralized manner among different spot markets with the objective of maximizing the total revenue is introduced in [35]. A market clearing pricing mechanism for VM exchanges between CPs is developed in [36], where a centralized broker dynamically adjusts a single VM price for the federation. Mihailescu et al. [37] relied on a centralized market broker in the federation, to which sellers publish resources and buyers send requests. Le et al. [38] addressed the cooperation problem among CPs but from the perspective of load balancing and electricity consumption. On the other hand, Lee et al. [39] addressed the cooperation problem from the perspective of providing a better quality of service by the CPs. The problem of resource cooperation amongst CPs when considering tasks with deadline constraints is presented in [40]. Finally, the work in [41] addresses the consumer's problem of selecting a VM reservation plan or request resources on demand with the objective of reducing the service cost. In general, existing approaches in the regulated setting are concerned only with the instantaneous CP gains. In this regard, Samaan [9] is the first to take into consideration the outcome of historical and future interactions among the CPs in making sharing decisions in the federation. Her model relies on a centralized federation broker that performs the VM allocations among the cloud providers. The author, through a combination of game theory and simulations, shows that her approach leads to a higher accumulated profit by the CPs. The model can also be extended to incorporate additional constraints for each CP to account for different issues (e.g., energy consumption). The policies proposed in [33] focus on helping providers to make the decisions on executing a job locally or remotely to improve providers' profit.

Differences with Our Work. Our work is both an extension as well as orthogonal to the papers mentioned above. It is an extension to the above mentioned works on decentralized control (with the exception of [6],[9]) because it models costs of borrowing cloud resources from peers, whereas the above mentioned works assume free resource sharing between SCs. The difference between [9] and our work is in the application setting. The former uses a spot cloud setting compared to our on-demand resource sharing setting. Our work is orthogonal to the above mentioned distributed approaches in that we study robust supply-demand markets for SCs, given an effective platform for enabling cooperation between cloud providers. In regard to works based on federated and centralized cloud sharing environments, our contribution is holistic compared to existing approaches that focus only on the single-dimensional problem of improving cost, QoS, etc. In comparison, we focus on the multi-dimensional problem of jointly satisfying the interests of different stakeholders in a market.

Markets for Cloud Cooperation. Buyya et al. [32] introduced a market oriented VM exchange model among clouds. Federations of hybrid clouds were also discussed in [33, 34]. Goiri et al. [33] proposed a novel federation management architecture and focused on building an accurate revenue function of CPs when taking one of the following decisions: (i) participate in the federation by outsourcing or insourcing capacity or (ii) turning off spare capacity. Similarly, Toosi et al. [34] provided a comprehensive analysis of the related costs and revenue associated with the decisions in (i) and (ii), of CPs in the federation. In [9], the author focussed on the problem of effective pricing and revenue maximization in incentivizing data centers to cooperate. *However, all of the above contributions focus only on market models that address the single-dimensional problem of improving cost, QoS, etc. In comparison, we focus on the multi-dimensional problem of jointly satisfying the interests of different stakeholders in a market.* In [27], the authors investigate a service exchange model for exchange networks that is generic and representative of many communication and distributed sytems (e.g., a cloud cooperative). As the main contribution, the authors in [27] fully characterize the market equilibria and relate them to the *max-min* fair market outcomes. *However, they do not (i) presume any type of money transfers (as in typical exchange economies), (ii) do not target reaching the social welfare optimal solution, which is atleast as good as the max-min fair solution and in most cases better, and (iii) do not address market instability issues that might arise due to resource volatility.*

7 Conclusion and Future Work

In this paper, we addressed the problem of effective resource sharing between small clouds (SCs). The problem arises mainly because of the increasing concerns amongst cloud customers of massive data centers about privacy breaches and data control by the latter. We modeled the problem as an efficient supply-demand market design task consisting of (i) autonomous SCs, (ii) their customers, and (iii) a regulator, as the market stakeholders. We showed that a centralized welfare allocation policy for the stakeholders by the regulator (in turn breaking SC autonomy) maximizes utilitarian social welfare at the static market equilibrium and results in the best/most efficient state at which the SDC markets could operate. Fortunately, courtesy *Arrow-Debreu* welfare theorems in welfare economics, this unique optimal operating point is also achieved by the autonomous SCs in price competition with one another, thereby guaranteeing no efficiency loss in a non-centralized market setting. However, the optimal market equilibrium point is prone to perturbations due to the dynamic nature of the SC market, thereby potentially leading to market disequilibrium. In this context, we designed a dynamic market mechanism based on Arrow and Hurwicz's disequilibrium process that uses the gradient play technique in game theory to converge upon the optimal static market efficient equilibrium from a disequilibrium state caused due to supply-demand perturbations, and results in market stability. We illustrated the stability and sustainability of dynamic SC markets and the high speed with which such markets converge to stable equilibria, through extensive experiments conducted on synthetic data.

As part of future work, we plan to design provably fast distributed algorithms to allow markets to roll back to efficient

equilibria when perturbed from an equilibrium state, and study dynamic SC markets under (i) a setting of imperfect competition between SCs, and (ii) a coalitional market setting where SCs have the capability to collude with one another.

REFERENCES

[1] H. Zhuang, R. Rahman, and K. Aberer, "Decentralizing the cloud: How can small data centers cooperate," in *IEEE International Conference on Peer-to-Peer Computing*, 2014.

[2] K. J. Arrow and L. Hurwicz, "On the stability of the competitive equilibrium i," *Econometrica*, vol. 26, no. 4, 1958.

[3] K. J. Arrow, H. D. Block, and L. Hurwicz, "On the stability of the competitive equilibrium ii," *Econometrica*, vol. 27, no. 1, 1959.

[4] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. H. Katz, A. Konwinski, G. Lee, D. A. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "A view of cloud computing," *Communications of the ACM*, vol. 53, 2010.

[5] P. D. Filippi and S. McCarthy, "Cloud computing: Legal issues in centralized architectures," in *International Conference on Internet, Law, and Politics*, 2011.

[6] Blippex, *Why we Moved Away from AWS*. http://blippex.github.io/updates/2013/09/23/why-we-moved-away-from-aws.htm.

[7] Mayer and Brown, *European and German Privacy Laws*. http://www.martindale.com.

[8] Wikipedia, *Business Cluster*. http://en.wikipedia.org.

[9] N. Samaan, "A novel economic sharing model in a federation of selfish cloud providers," *IEEE Transactions on Parallel and Distributed Systems*, vol. 25, no. 1, 2014.

[10] B. Rochwerger, D. Breitgand, A. Epstein, D. Hadas, I. Loy, K. Nagin, J. Tordsson, C. Ragusa, M. Villari, S. Clayman, E. Levy, A. Maraschini, P. Massonet, H. Munoz, and G. Toffetti, "Reservoir: When one cloud is not enough," *Computer*, vol. 44, no. 3, 2011.

[11] B. Sotomayor, R. S. Montero, I. M. Llorente, and I. Foster, "Virtual infrastructure management in private and hybrid clouds," *IEEE Internet Computing*, vol. 13, no. 5, 2009.

[12] A. Mas-Collel, M. D. Whinston, and J. R. Green, *Microeconomic Theory*. Oxford University Press, 1995.

[13] J. S. Shamma and G. Arslan, "Dynamic fictitious play, dynamic gradient play, and distributed convergence to nash equilibria," *IEEE Transactions on Automatic Control*, vol. 50, no. 3, 2005.

[14] D.Fudenberg and J.Tirole, *Game Theory*. MIT Press, 1991.

[15] S. D. Flam, "Equilibrium, evolutionary stability, and gradient dynamics," *International Game Theory Review*, vol. 4, no. 4, 2002.

[16] G. Ananthnarayanan, M. Hung, X. Ren, I. Stoica, A. Wierman, and M. Yu, "Grass: Trimming stragglers in approximation analytics," in *NSDI*, 2014.

[17] O. Candogan, K. Bimpikis, and A. Ozdaglar, "Optimal pricing in networks with externalities," *Operations Research*, vol. 60, no. 4, 2012.

[18] J.-J. Laffont and J. Tirole, *A Theory of Incentives in Procurement and Regulation*. MIT Press.

[19] J. Bulow, J. Geneakolos, and P. Klemperer, "Multi-market oligopoly: Strategic substitutes and complements," *Journal of Political Economy*, vol. 93, 1985.

[20] J. J. Laffont and J.Tirole, *A Theory of Incentives in Procurement and Regulation*. MIT Press, 1993.

[21] S.Boyd and L.Vanderberghe, *Convex Optimization*. Cambridge University Press, 2005.

[22] Y. Nesterov and A. Nemirovski, *Interior-Point Polynomial Algorithms in Convex Programming*. SIAM Studies in Applied Mathematics, 1994.

[23] A. Bacciotti and L. Rosier, *Liapunov Functions and Stability in Control Theory*. Springer, 2005.

[24] D. P. Bertsekas and J. Tsitsiklis, *Parallel and Distributed Computation*. Prentice-Hall, 1992.

[25] A. Chandra, J. Weismann, and B. Heintz, "Decentralized edge clouds," *IEEE Internet Computing*, vol. 17, no. 5, 2013.

[26] O. Babaoglu, M. Marzolla, and M. Tamburini, "Design and implementation of a p2p cloud system," in *ACM Symposium on Applied Computing*, 2012.

[27] L. Georgiadis, G. Iosifidis, and L. Tassiulas, "Exchange of services in networks: Competition, cooperation, and fairness," in *ACM SIGMETRICS*, 2015.

[28] H. Wang, F. Wang, J. Liu, and J. Green, "Measurement and utilization of customer-provided resources for cloud computing," in *IEEE INFOCOM*, 2012.

[29] B. Rochwerger, D. Breitgand, A. Epstein, D. Hadas, I. Loy, K. Nagin, J. Tordsson, C. Ragusa, M. Villari, S. Clayman, E. Levy, A. Maraschini, P. Massonet, H. Munoz, and G. Tofetti, "Reservoir - when one cloud is not enough," *Computer*, vol. 44, no. 3, 2011.

[30] B. Rochwerger, E. Levy, A. Galis, K. Nagin, I. M. Llorente, R. Montero, Y. Wolfsthal, E. Elmroth, J. Caceres, M. Ben-Yehuda, W. Emmerich, and F. Galan, "The reservoir model and architecture for open federated cloud computing," *IBM Journal of Research and Development*, vol. 53, no. 4, 2009.

[31] A. Avetisyan, "Open cirrus: A global cloud computing testbed," *Computer*, vol. 43, no. 4, 2010.

[32] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud computing and emerging it platforms: Vision, hype, and reality for delivering computing as the fifth utility," *Future Generation Computer Systems*, vol. 25, no. 6, 2009.

[33] I. Goiri, J. Guitart, and J. Torres, "Economic model of a cloud provider operating in a federated cloud," *Information Systems Frontiers*, vol. 12, 2012.

[34] A. Toosi, R. Calheiros, R. Thulasiram, and R. Buyya, "Resource provisioning policies to increase iaas provider's profit in a federated cloud environment," in *IEEE HPCC*, 2011.

[35] Q. Zhang, E. Gurses, R. Boutaba, and J. Xiao, "Dynamic resource allocation for spot markets in clouds," in *USENIX Hot-ICE*, 2011.

[36] E. Gomes, Q. Vo, and R. Kowalczyk, "Pure exchange markets for resource sharing in federated clouds," *Concurrency and Computation: Practice and Experience*, vol. 24, no. 9, 2012.

[37] M. Mihailescu and Y. Teo, "Dynamic resource pricing on federated clouds," in *IEEE/ACM CCGrid*, 2010.

[38] K. Le, R. Bianchini, J. Zhang, Y. Jaluria, J. Meng, and T. Nguyen, "Reducing electricity cost through virtual machine placement in high performance clouds," in *SC*, 2011.

[39] Y. C. Lee, C. Wang, A. Y. Zomaya, and B. B. Zhou, "Profit-driven service request scheduling in clouds," in *IEEE/ACM CCGrid*, 2010.

[40] R. V. den Bossche, K. Vanmechelen, and J. Broeckhove, "Cost-optimal scheduling in hybrid iaas clouds for deadline constrained workloads," in *IEEE CLOUD*, 2010.

[41] S. Chaisiri, B. Lee, and D. Niyato, "Optimization of resource provisioning cost in cloud computing," *IEEE Transactions on Services and Computing*, vol. 5, no. 2, 2012.

[42] G. Wang, M. Negrete, A. Kowli, E. Shafieepoorfard, S. Meyn, and U. V. Shanbhag, *Dynamic Competitive Equilibria in Electricity Markets*. Springer-Verlag, 2011.

APPENDIX

In this section we provide the proofs for Theorems 1-3.

Proof of Theorem 1. The equilibrium (x_1^*, x_2^*) when setting κ_j^1, κ_j^2 to zero, is a solution of the following.

$$\rho_i^* - \beta_i^r vm_i^{r*} - \alpha_i^r = 0, \forall i \in SDC. \tag{23a}$$

$$\rho_i^* - \beta_i^b vm_i^{b*} - \alpha_i^b = 0, \forall i \in SDC. \tag{23b}$$

$$\rho_i^* - \beta_i^{pc} vm_i^{pc*} - \alpha_i^{pc} = 0, \forall i \in SDC. \tag{23c}$$

$$\beta_i^{type} vm_i^{type} + \alpha_i^{type} - \rho_i^* = 0, \forall i \in C, type \in \{e, c, s, ag\}. \tag{23d}$$

$$\sum_{j \in C_i} vm_j^{ag}(1 - \kappa_j^1 - \kappa_j^2) = (vm_i^r + vm_i^b + vm_i^{pc}), \forall i \in SDC. \tag{23e}$$

Using Theorem 3 in [42], strong duality implies that equilibrium $(x_1^*, x_2^*$ exists is identical to the solution of the KKT conditions in (10a)-(10e). It can be seen that (25a) follows by replacing the cost function for SDCs in (2)-(4) in (10a). Similarly, (25b) follows by replacing the utility function of SDC customers in (5)-(8) in (10d). Furthermore (25c) is identical to (10e). Thus, $(x_1^*, x_2^*$ is identical to the equilibrium in (10a)-(10e). Thus, we proved Theorem 1. ■.

Proof of Theorem 2. Since strong duality holds, it follows from Theorem 1 that equilibrium $(x_1^*, x_2^* \in E$ exists. We first prove the stability of this equilibrium point and then proceed to its asymptotic stability. Differentiating the positive definite Lyapunov function $V(y_1, y_2) = y_1^T P_1 y_1 + y_2^T P_2 y_2$, with respect to time where $y_1 = x_1 - x_1^*$ and $y_2 = x_2 - x_2^*$, and by using the non-expansive property of the projection operation, we have

$$V(\dot{y_1}, y_2) \leq y_1^T (P_1 A_1 + A_1^T P_1) y_1 + y_1^T P_1 A_2 y_2 + y_2 A_2^T P_1 y_1 \tag{24}$$

If A_1 is Hurwitz, for any $Q > 0$, there exists a positive definite matrix P_1 such that $P_1 A_1 + A_1^T P_1 = -Q$. Let $\lambda_{\min}(Q)$ denote the minimum eigenvalue of Q. Since P_2 is a symmetric positive definite matrix with a set n orthogonal, real, and non-zero eigenvectors $x_1, ..., x_n$, can be written as

$$P_2 = \sum_{i=1}^{n} \lambda_i x_i x_i^T,$$

where $\lambda_i > 0$ is the eigenvalue corresponding to x_i. We can expand the vector $VM^{\max}$ using the orthogonal vector w_i as

$$VM^{max^T}[1]_{n \times n} P_2 y_2 \geq \lambda_{\min}(P_2) \psi_{\min} ||y_2||_2, \tag{25}$$

where $\psi_{\min} = \min(\psi_i), \forall i = 1, ..., n$. Now let

$$\beta \geq ||P_1 A_2 + R^T[1]_{n \times n} P_2||_2.$$

Using (26) and (27), we obtain

$$\begin{aligned} \dot{V(y_1, y_2)} &\leq -\lambda(Q)\left(||y_1||_2 - \frac{\beta}{\lambda_{\min}(Q)}||y_2||_2\right)^2 \\ &- ||y_2||\left(2\lambda_{\min}(P_2\psi_{\min} - \frac{\beta^2}{\lambda_{\min}(Q)}||y_2||\right). \end{aligned}$$

For all $\Omega_{\max} \subsetneq D$, it follows that for all solutions beginning in Ω_{max}, $V \leq 0$. Hence, the equilibrium is stable and $\Omega_{\max}$ is the region of attraction.

Since the initial conditions start in Ω_Δ and the latter is a strict subset of D_Δ, y_2 cannot be equal to $2\lambda_{\min}(P_2)\psi_{\min}\frac{\lambda_{\min}(Q)}{\beta^2}$. This in turn implies that $(||y_1||, ||y_2|| = (0,0)$ is the only invariant set. Hence, all solutions starting in Ω_Δ converge to the equilibrium point $(x_1, x_2) = (x_1^*, x_2^*)$. Thus, we proved Theorem 2. ■

Proof of Theorem 3. Differentiating the Lyapunov function $V(y_1, y_2)$ along the trajectories of (16), we get

$$\dot{V(y_1, y_2)} \leq -a_\Delta\left(||y_1|| - \frac{\beta}{a_\Delta}||y_2||\right)^2 - ||y_2||\left(e - \frac{\beta^2}{a_\Delta}||y_2||\right), \tag{26}$$

where $a_\Delta = \lambda_{\min}(Q) - 2||P_1||\pi_{SDC} + 2||P_1||\pi_C$, and $e = 2\lambda_{\min}(P_2)\psi_{\min}$.

From (24) it follows that $a_\Delta > 0$. Therefore, (25) implies that for all $\Omega_{c_{max}} \subsetneq D_\Delta$, for all solutions beginning in Ω_Δ, $\dot{\bar{V}} \leq 0$. Hence, the market equilibrium state is stable, and Ω_Δ is the region of attraction.

The asymptotic stability of the perturbed market can be shown via the following argument: since the initial conditions start in Ω_Δ and the latter is a strict subset of D_Δ, y_2 cannot be equal to $2\lambda_{\min}(P_2)\psi_{\min}\frac{\lambda_{\min}(Q)}{\beta^2}$. This in turn implies that $(||y_1||, ||y_2|| = (0,0)$ is the only invariant set. Hence, all solutions starting in Ω_Δ converge to the equilibrium point $(x_1, x_2) = (x_1^*, x_2^*)$. Thus, we have proved Theorem 3. ■

Ranjan Pal is a Research Scientist at the University of Southern California (USC), affiliated with both the Electrical Engineering and Computer Science departments, where he co-leads the Quantitative Evaluation and Design Group (QED). He received his PhD in Computer Science from USC in 2014, and was the recipient of the Provost Fellowship throughout his PhD studies. During his PhD, Ranjan held visiting scholar positions at the School of Engineering and Applied Science, Princeton University, USA, and Deutsch Telekom Research Laboratories (T-Labs), Berlin, Germany. His primary research interests lie in the performance modeling, analysis, and design of cyber-security, privacy, communication networks, and the Smart Grid, using tools from economics, game theory, applied probability, algorithms, information theory, and mathematical optimization. His research on cyber-insurance has generated press interests from USC News, and MIT Technology Review. Ranjan has also consulted for Accel Partners on cyber-insurance. He is a member of the IEEE and the ACM.

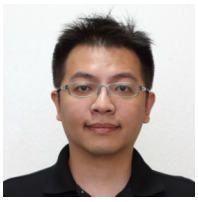

Sung-Han Lin is currently a Ph.D. Candidate in the Department of Computer Science, University of Southern California (USC). He received his M.S and B.S. degree in Computer Science and Information Engineering from National Taiwan University in 2008 and 2006 respectively. His main research interests lie in the performance modeling, analysis, and design of large-scale distributed systems, including cloud computing systems, peer-to-peer networking, and multimedia systems. Sung-Han is a member of the IEEE and the ACM.

Leana Golubchik is a Professor in the Computer Science and EE-Systems Departments at USC. Prior to that she was an Assistant Professor and then an Associate Professor at the University of Maryland at College Park and an Assistant Professor at Columbia University. Leana received her PhD from UCLA in 1995. She is a past Chair (2003 - 2005) and Vice Chair (2001 - 2003) of ACM SIGMETRICS (2003 - 2005) and a member of its Board of Directors (1999 - 2001, 2005 - 2007). She was a guest co-editor for special issues of IEEE TKDE, the Parallel Computing Journal, and the International Journal of Intelligent Systems, a program co-chair of the 2001 Joint ACM SIGMETRICS/Performance Conference and MIS99, and a program committee member of numerous conferences. Leana received several awards, including the NSF CAREER award, the Okawa Foundation award, and the IBM and NSF Doctoral Fellowships. She is a member of the IFIP WG 7.3 and Tau Beta Pi.